Load Balancing for Adaptively Refined Grids

Gerhard Zumbusch *

The solution of partial differential equations on a parallel computer is usually done by a data parallel approach. The grid is partitioned and mapped onto the processors. However, partitioning of unstructured meshes and adaptively refined meshes in general is an NP-hard problem and heuristics are needed. In this paper a parallelisable and cheap method based on space-filling curves is analysed. Quasi-optimal estimates are derived for partitions of adaptively refined grids.

MSC2000 classification: 65n50, 90c35, 68w10, 68q25, 42c40

1 Introduction

When it comes to the parallelisation of a PDE solver, usually a data parallel approach is appropriate [5,8,9]. We will discuss the grid partitioning or load balancing problem. We will do this by means of space-filling curves [12]. While the general graph partitioning problem [11] is NP-hard [1], grids provide more information, namely the location of the nodes in space. Hence it seems to be more likely to construct grid partitions of acceptable quality in polynomial time than for the pure graph partitioning problem. Nevertheless, we are looking for a heuristic which is linear in time, like e.g. a linear complexity multigrid solver, and additionally which can be parallelised efficiently. This can indeed be done by space-filling curves, which enumerate all elements or nodes of a grid, see [10]. The partitioning algorithm simply chooses a complete interval of numbers, which translates into a geometric domain.

The main focus of this paper is on the analysis of this heuristic. Basically we are able to prove that the partitions are quasi-optimal. This means that the cut size of the graph is bounded by the optimal cut size times a constant, independently of the mesh size. This is true both for uniformly refined grids and for certain types of adaptively refined grids. Although some advanced graph partitioners may produce slightly better cut sizes in numerical experiments than space-filling curves, many come without any proof on the quality of the partitions generated. In [13] the quasi-optimal quality of space-filling curve partitions of uniform grids was proven. We will generalise these results to the case of adaptively refined grids.

*Institut für Angewandte Mathematik, Universität Bonn, Wegelerstr. 6, D-53115 Bonn, Germany, e-mail zumbusch@iam.uni-bonn.de, URL http://wissrech.iam.uni-bonn.de/people/zumbusch.html

2 Generalised Sierpiński Curves

The Sierpiński space-filling curve was defined originally on a triangle, see [12]. However, the curve can be generalised to arbitrary unstructured triangulations. Given a sequence of meshes of polyhedral elements which is created by element refinement and which element sizes all tend to zero. A numbering of the elements of the coarsest mesh gives an initial curve. A finite set of refinement rules is used to subsequently construct finer meshes and the curve is refined accordingly. Note that a generalised Sierpiński curve is space-filling in the limit case of zero element sizes. We will use the generalised Sierpiński curves for grid partitioning.

Furthermore, we define a quasi-uniform refineable family T of grids as a nested sequence of grids τ_i which covers a domain Ω. Each grid consists of a finite number of non-overlapping polygonally shaped elements $e_{i,j} \in \tau_i$ by the conditions:

(a) $\overline{\bigcup_j e_{i,j}} = \Omega$ and $e_{i,j} \cap e_{i,k} = \emptyset$ for $j \neq k$, (grid condition),

(b) the element size has uniformly limit zero, $\lim_{i\to\infty} \max_j \mathrm{vol}(e_{i,j}) = 0$ (T is dense),

(c) either $e_{i,j} \in \tau_{i+1}$ or a set K of smaller elements exists with $\overline{e_{i,j}} = \overline{\bigcup_{k\in K} e_{i+1,k}}$ (nestedness),

(d) the cardinality $|K| \leq c_c$ is bounded from below independent of i, j (finite number of children),

(e) the ratio of volumes $\mathrm{vol}(e_{i+1,k})/\mathrm{vol}(e_{i,j}) \geq c_{\min} > 0$ for $k \in K$ is bounded from below independent of i, j (local volume condition),

(f) all interior angles of the element $e_{i,j}$ are bounded from below $\geq \alpha_0 > 0$ independent of i, j (angle condition),

(g) the number of vertices of the element $e_{i,j}$ is bounded by $\leq c_v$ (vertex condition).

Many grids used in the FEM will be quasi-uniform refineable, if we consider the limit case of step size $h \to 0$: Of course uniformly refined grids, starting on the unit square or an arbitrary unstructured mesh fall into this category. Furthermore, grids created by adaptive refinement by quad-tree or oct-tree refinement, with closure rules and any stable bisection schemes for triangle and tetrahedron grids, or hybrid grids.

A partition A of a grid contains a number of nodes, called volume v and has a certain number of boundary faces, called surface s. The associated graph has v vertices and the removal of A cuts s edges of the graph. The isoperimetric number is the maximum ratio v/s. It is of the order $\mathcal{O}(n^{1/d})$ for a domain in d dimensions and n vertices for graphs we are interested in. However, it is NP hard to find the optimum or just an approximation $n^{1/d-\epsilon}$ away [1]. Hence we call a result of $s/v = \mathcal{O}(n^{1/d})$ quasi-optimal for partitions of the desired size v. This was proven for uniform grids in [13]. Furthermore it is known for certain types of spectral partitioning methods of α-overlap graphs, which are usually not of arbitrary size v [7].

3 β-adaptive Grid Refinement

We define an β-adaptive family T of grids as a quasi-uniform refineable family of grids with an additional condition. Each element $e_{i,j}$ of a grid τ_i is subdivided into at least two smaller elements in grid τ_{i+1} reflected by the constants $1 < C_2 \leq C_1 < \infty$ and $C_2 \leq \left|\left\{e_{i+1,k} \in \tau_{i+1} | e_{i+1,k} \subset e_{i,j}\right\}\right| \leq C_1$. The value $\beta := \log C_1 / \log C_2 \geq 1$ is a measure for the differences of the number of refinement steps taken in different areas of the domain Ω.

L e m m a 1. *Let $f : I \mapsto \Omega$ be a generalised Sierpiński curve on a β-adaptive family of grids. Then $s \leq C_{\text{part}} C_1 \cdot v^{\beta \frac{d-1}{d}}$ holds for a partition with v the number of elements and s the number of element faces.*

Many grids used for the solution of partial differential equations are in fact β-adaptive with some value β. Of course sequences of uniform and quasi-uniform grids fulfil the conditions. Furthermore, adaptive grid refinement procedures may lead to β-adaptive grids if the grid refinement is controlled in a certain way. However, there are also grids with $\beta \geq \frac{d}{d-1}$ such that the estimate is not useful.

4 γ-adaptive Grid Refinement

Under some sharper assumptions than β-adaptivity where we restrict the number of elements by subdivision, we can give a tighter estimate. In order to do this, we generalise the definition of β-adaptivity. We define the function $\gamma \geq 0$ for a quasi-uniform refineable family of grids by the following condition. Each element $e_{i,j}$ of a grid τ_i is subdivided into two or more smaller elements in grid τ_{i+1} reflected by the bounds $\underline{\gamma}(1)$ and $\overline{\gamma}(1)$ with $1 < \underline{\gamma}(1) \leq \overline{\gamma}(1) < \infty$. We assume that uniform grid refinement is associated to the lower bound. Furthermore, this number of elements is refined uniformly in any case, such that only up to $\overline{\gamma}(1) - \underline{\gamma}(1)$ elements are created additionally by another procedure.

The total number of elements created in refinement step k on grid τ_{i+k} is a number of elements between the bounds $\underline{\gamma}(k)$ and $\overline{\gamma}(k)$ with $1 < \underline{\gamma}(k) \leq \overline{\gamma}(k) < \infty$ with monotonically increasing $\underline{\gamma}$. Again, the uniform refinement part is at least $\underline{\gamma}(k)\overline{\gamma}(k-1)/\underline{\gamma}(k-1)$ with an additional part of elements created by another procedure. For technical reasons, we assume that $\underline{\gamma}(k)$ increases exponentially, i.e. $\underline{\gamma}(k) = C_\gamma^k$ with $C_\gamma > 1$. This can be achieved by a reparametrisation of the functions γ and the grids τ, if a single set of uniform grid refinement rules is used.

The definition does not restrict the set of quasi-uniform refineable grids like β-adaptivity. However, we are interested in families of grids with small ratios $\overline{\gamma}/\underline{\gamma}$. For example for β-adaptive grids we immediately have the constant functions $\underline{\gamma}(i) = C_2^i$ and $\overline{\gamma}(i) = C_2^i$. Conversely, a γ-adaptive grid will also be β-adaptive if $\beta := \overline{\lim}_{i \to \infty} \frac{\log \overline{\gamma}(i)}{\log \underline{\gamma}(i)}$ is finite, i.e. the ratio $\overline{\gamma}/\underline{\gamma}$ is bounded.

In more generality, the number of elements created by uniform refinement can be at least some function C^k with $C < C_\gamma$. In such a case, the following estimates will be weaker by the exponent $\log C_\gamma / \log C$, analogously to the estimates on β-adaptive grids with $\beta > 1$.

L e m m a 2. *Let $f : I \mapsto \Omega$ be a self-similar $1/d$-Hölder continuous space-filling curve parametrised by volume, on a γ-adaptive family of grids with $\sum_{j=0}^{\infty}(\overline{\gamma}(j) - C_\gamma \overline{\gamma}(j -$*

$1))/C_\gamma^{\frac{j(d-1)}{d}} < \infty$. *Then the estimate* $s \leq C_{\text{part}} \cdot v^{\frac{d-1}{d}}$ *holds for a partition with* v *the number of elements and* s *the number of element faces. However, if the sum grows linearly in* j*, we obtain the weaker estimate of* $s \leq C_{\text{part}} \cdot v^{\frac{d-1}{d}} \log v$.

5 Adaptively Refined Grid Example

We will discuss some typical situations of singularities and their related functions $\overline{\gamma}$. Now the strength α of an asymptotic r^α-expansion for singularities due to boundary conditions does not play the prominent role as for β-adaptive grids, but the relative area of grid refinement. Note that the case of codimension one grid refinement it is easy to see that the surface s is proportional ($d > 2$) or only logarithmically ($d = 2$) lower than the volume v.

We construct a quasi-uniform refinable sequence of grids. Starting with an arbitrary initial grid τ_0, each grid is created by a uniform refinement step and an adaptive refinement step. The mesh parameter h is divided by two at the uniform refinement step and the minimal mesh parameter is divided by two again at the adaptive refinement step. Elements adjacent to a g-dimensional sub-manifold S of Ω are refined during the adaptive refinement step. We obtain a γ-adaptive grid with $C_\gamma = 2^d$ and a function $\gamma(j) = \mathcal{O}(4^{gj})$. With lemma 2 follows that the surface to volume ratio of generalised Sierpiński space-filling curve partitions is bounded for $g < (d-1)/2$ and is logarithmically bounded for $g = (d-1)/2$. Further results are assembled in the following table. Entries marked by an asterisk * can be computed directly for this example.

	$d=2$	$d=3$	$d=4$	$d=5$	$d=6$
$g=0$	$\leq \sqrt{v}$	$\leq v^{2/3}$	$\leq v^{3/4}$	$\leq v^{4/5}$	$\leq v^{5/6}$
$g=1$	$\geq v/\log v$	$\leq v^{2/3}\log v$	$\leq v^{3/4}$	$\leq v^{4/5}$	$\leq v^{5/6}$
$g=2$		$= v$	$= v^*$	$\leq v^{4/5}\log v$	$\leq v^{5/6}$
$g=3$			$= v$	$= v^*$	$= v^*$
$g=4$				$= v$	$= v^*$
$g=5$					$= v$

6 Best n-Term Approximation Grids

A systematic way to construct adaptively refined grids is by best n-term approximation. The approximation error is minimised under the restriction that only n degrees of freedom may be used. The idea is to choose these degrees in an optimal way. This is a non-linear approximation procedure. Technically, this best approximation as the solution of a discrete optimisation problem is difficult to achieve. However, there are possibilities to find approximations close to the best approximation in the following way: Given a stable Riesz basis, the size of the weighted coefficients related to an expansion in the stable basis is an indicator of the importance of the basis function. Hence, it is sufficient to choose the n largest weighted coefficients for an approximation, which is up to the stability constants the degrees-of-freedom to error relation of the best n-term approximation [3]. We define the best approximation error as $\sigma_{n,t}(g) = \inf_{|\Lambda| \leq n,\ d_\lambda \in \mathbb{R}} \left\| g - \sum_{\lambda \in \Lambda} d_\lambda \psi_\lambda \right\|_{H^t(\Omega)}$ for a set of basis functions

ψ_λ and weights d_λ. For the solution of a Poisson problem with right hand side $f \in B^{\alpha-1}_{L_2(\Omega)}$, $\alpha \geq 1$ and a stable wavelet basis with at least α vanishing moments, we obtain for the procedure choosing the n largest weighted coefficients d_λ of $\sum_{n=1}^{\infty} \left(n^{s/d}\sigma_{n,1}(u)\right)^{\frac{s-1}{d}+\frac{1}{2}} < \infty$ for all $0 < s < \min(\frac{d}{2(d-1)}, \frac{\alpha+1}{3})$, see [2]. This is the optimal degrees-of-freedom to error rate $n^{-1/d}$ of uniform grid refinement and H^2 regular problems, as long as $\alpha \geq \frac{d+2}{2d-2}$.

The best n-term approximation gives an explicit characterisation of (quasi-) optimal grids. We can study singularities which may occur for the solution of the Poisson equation. Basically, functions of type $r^{-\alpha}$ with $\alpha > 0$ are of interest, which model the behaviour of corner singularities due to the geometry of the domain Ω and due to boundary conditions. The radius r is given in polar coordinates for point singularities or e.g. in cyclindrical coordinates for edge singularities. The wavelet coefficients of a representation of the function can be computed explicitly, such that the best n-term grids are constructed in the following. The convergence rate in L_2 norm is of optimal order $n^{-1/d}$.

We assume a dyadic grid refinement and characterise the grid level by level for a point singularity and $d = 1$. Higher dimensions and other types of refinement can be obtained by direct products of the one-dimensional case and products with uniform refinement. For a given threshold δ, all nodes of a dyadic level $\underline{l}$ and coarser are present. Furthermore, there is a maximum dyadic level $\overline{l}$ and on a dyadic level l in between, $n_l = C\delta^{-a}b^l$ nodes are present, with constants $C, a, b > 0$ and $b < 1$. On each dyadic level only the nodes closest to the singularity are present. Hence the bounds on the levels can be calculated by $n_{\overline{l}} = 1$ and $n_{\underline{l}} = 2^{\underline{l}} + 1$. We parametrise the grids by a sequence of thresholds, such that we obtain a γ-adaptive family of grids. For a lower bound of $C_\gamma = 2^d$ we need a sequence of $\underline{l} \in \mathbb{N}$, which results in a geometrically decreasing δ by a factor $(b/2)^{1/a}$. Each grid τ_i with $\underline{l} = i$ is geometrically graded. The next finer grid τ_{i+1} can be obtained by a refinement of τ_i. Each element is refined uniformly. Furthermore, a constant number c of elements is created additionally next to the singularity on levels $l > \overline{l}(\tau_i)$. Hence $\overline{\gamma}(i+1) = C_\gamma\overline{\gamma}(i) + c$ and with lemma 2 optimal space-filing curve partitions can be obtained.

For the adaptive grid refinement toward a g dimensional manifold like for edge singularities ($g = 1$), the upper bound $\overline{\gamma}$ changes. The property $\overline{l} = \underline{l}(1 - \frac{\log 2}{\log b})$ with $b < 1$ holds. Hence more elements are created in a step from τ_i to τ_{i+1} in addition to the uniform refinement. We obtain a g dependent part of the upper bound as

$$\overline{\gamma}(i) = \mathcal{O}\left(2^{g\bar{l}(i)} + \underline{\gamma}(i)\right) = \mathcal{O}\left(2^{ig(1-\frac{\log 2}{\log b})} + 2^{id}\right).$$

Inserting this into lemma 2, optimal space-filing curve partitions are guaranteed for $g < (d-1)/(1 - \frac{\log 2}{\log b})$. The constant b can be computed explicitly for each wavelet type and error norm. This leads to an analog table, which depends on α and the number of vanishing moments of the wavelet.

7 Anisotropic Curves

Note that the Besov regularity for singularities with $g > 0$ of $r^{-\alpha}$ type is much lower than for the point singularity $g = 0$. Hence best n-term approximation does not lead to an $\mathcal{O}(n^{-1/d})$ convergence. This can also explained by the geometric nature of the singularity, where

anisotropic shape functions are more efficient. The approximation results can be improved by anisotropic grids. Furthermore, we have seen that the space-filling curve partitions of adaptively refined grids can give insufficient locality in this case. If g is too large compared to d, i.e. $g = d - 1$, the standard space-filling curves will lead to partitions with $s = \mathcal{O}(v)$. Analytically, the domain partitioning problem is relatively easy in such cases: Cut the domain into g-dimensional slices.

The discrete curve, which we will call an 'anisotropic' space-filling curve, is constructed by a direct product of lower-dimensional space-filling curves. This does not lead to a continuous space-filling curve in the limit case. However, the partitions defined by the curve serve our purpose. We can use the locality estimates on the curve orthogonal to the slices to demonstrate the locality of the partitions with respect to the adaptive grid refinement. In the case of $g = d - 1$ dimensional refinement, i.e. codimension one, the surface of the slices is of the order of adaptive refinement with $g = d - 2$, e.g. sufficient locality for $d = 2$ with $s \leq \mathcal{O}(\sqrt{v})$ and for $d = 3$ with $s \leq \mathcal{O}(v^{2/3} \log v)$.

8 References

1 Bui, T. N., Jones, C.: Finding good approximate vertex and edge partitions is NP-hard. Inf. Process. Letters **42** (1992), 153–15.

2 S. Dahlke, S., Dahmen, W., DeVore, R. Nonlinear approximation and adaptive techniques for solving elliptic operator equations. in: Multiscale Techniques for PDEs, Dahmen, W., Kurdilla, A., Oswald, P., Academic Press (1997), 237–284.

3 DeVore, R.: Nonlinear Approximation. Acta Numerica **7** (1998), 51–150.

4 Gotsman, C., Lindenbaum, M.: On the metric properties of discrete space-filling curves. IEEE Trans. on Image Processing **5** (1996), 794–797.

5 Griebel, M., Zumbusch, G.: Hash-Storage Techniques for Adaptive Multilevel Solvers and their Domain Decomposition Parallelization. in: Proc. Domain Decomposition Methods 10. Mandel, J., Farhat, C., Cai, X.-C., Contemporary Mathematics **218**, AMS (1998), 279–286.

6 Griebel, M., Zumbusch, G.: Parallel adaptive subspace correction schemes with applications to elasticity. Computer methods in applied mechanics and engineering. **184** (2000), 303–332.

7 Miller, G. L., Teng, S-H., Thurston, W., Vavasis, S. A.: Geometric separators for Finite Element Meshes. SIAM J. Sci. Comput. **10** (1998), 364–386.

8 Oden, J. T., Patra, A., Feng, Y.: Domain Decomposition for Adaptive hp Finite Element Methods. in: Proc. Domain Decomposition 7, Contemporary Mathematics **180**, AMS (1994), 295–301.

9 Parashar, M., Browne, J. C.: On partitioning dynamic adaptive grid hierarchies. in: Proc. 29th Annual Hawai Int. Conf. System Sciences (1996), 604–613.

10 Pilkington, J. R., Baden, S. B.: Partitioning with Spacefilling Curves. Report CS94-349, UCSD, Dept. Computer Science (1994).

11 Pothen, A.: Graph Partitioning Algorithms with Applications to Scientific Computing. in: Parallel Numerical Algorithms, Keyes, D. E., Sameh, A., Venkatakrishnan, V. (eds.), Kluwer (1997), 323–368.

12 Sagan, H.: Space-Filling Curves, Springer (1994).

13 Zumbusch, G.: On the Quality of Space-filling Curve Induced Partitions. Z. Angew. Math. Mech. **81** (2001), Suppl. 1, 25–28.